arXiv:1608.06403v1 [cs.GT] 23 Aug 2016

Phased Exploration with Greedy Exploitation in Stochastic Combinatorial Partial Monitoring Games

Sougata Chaudhur
Department of Statistics
University of Michigan Ann Arbor
sougata@umich.edu

Ambuj Tewari
Department of Statistics and EECS
University of Michigan Ann Arbor
tewaria@umich.edu

Abstract

Partial monitoring games are repeated games where the learner receives feedback that might be different from adversary's move or even the reward gained by the learner. Recently, a general model of combinatorial partial monitoring (CPM) games was proposed [1], where the learner's action space can be exponentially large and adversary samples its moves from a bounded, continuous space, according to a fixed distribution. The paper gave a confidence bound based algorithm (GCB) that achieves $O(T^{2/3}\log T)$ distribution independent and $O(\log T)$ distribution dependent regret bounds. The implementation of their algorithm depends on two separate offline oracles and the distribution dependent regret additionally requires existence of a unique optimal action for the learner. Adopting their CPM model, our first contribution is a Phased Exploration with Greedy Exploitation (PEGE) algorithmic framework for the problem. Different algorithms within the framework achieve $O(T^{2/3}\sqrt{\log T})$ distribution independent and $O(\log^2 T)$ distribution dependent regret respectively. Crucially, our framework needs only the simpler "argmax" oracle from GCB and the distribution dependent regret does not require existence of a unique optimal action. Our second contribution is another algorithm, PEGE2, which combines gap estimation with a PEGE algorithm, to achieve an $O(\log T)$ regret bound, matching the GCB guarantee but removing the dependence on size of the learner's action space. However, like GCB, PEGE2 requires access to both offline oracles and the existence of a unique optimal action. Finally, we discuss how our algorithm can be efficiently applied to a CPM problem of practical interest: namely, online ranking with feedback at the top.

1 Introduction

Partial monitoring (PM) games are repeated games played between a learner and an adversary over discrete time points. At every time point, the learner and adversary each simultaneously select an action, from their respective action sets, and the learner gains a reward, which is a function of the two actions. In PM games, the learner receives limited feedback, which might neither be adversary's move (full information games) nor the reward gained (bandit games). In *stochastic* PM games, adversary generates actions which are independent and identically distributed according to a distribution fixed before the start of the game and unknown to the learner. The learner's objective is to develop a learning strategy that incurs low regret over time, based on the feedback received during the course of the game. Regret is defined as the difference between cumulative reward of the learner's strategy and the best fixed learner's action in hindsight. The usual learning strategies in online games combine some form of exploration (getting feedback on certain learner's actions) and exploitation (playing the perceived optimal action based on current estimates).

29th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

There has been a substantial advance in our understanding of *finite* PM games. Starting with early work in the 2000s [2, 3], this body of research reached a culmination point with a comprehensive and complete classification of finite PM games [4]. We refer the reader to these works for more references and also note that newer results continue to appear [5]. Finite PM games restrict both the learner's and adversary's action spaces to be finite, with a very general feedback model. All finite partial monitoring games can be classified into one of four categories, with minimax regret $\Theta(T)$, $\Theta(T^{2/3})$, $\Theta(T^{1/2})$ and $\Theta(1)$. The classification is governed by *global* and *local observability* properties pertaining to a game [4]. Another line of work has extended traditional multi-armed bandit problem (MAB) [6] to include combinatorial action spaces for learner (CMAB) [7, 8]. The combinatorial action space can be exponentially large, rendering traditional MAB algorithms designed for small finite action spaces, impractical with regret bounds scaling with size of action space. The CMAB algorithms exploit a finite subset of base actions, which are specific to the structure of problem at hand, leading to practical algorithms and regret bounds that do not scale with, or scale very mildly with, the size of the learner's action space.

While finite PM and CMAB problems have witnessed a lot of activity, there is only one paper [1] on combinatorial partial monitoring (CPM) games, to the best of our knowledge. There the authors combined the combinatorial aspect of CMAB with the limited feedback aspect of finite PM games, to develop a CPM model. The model extended PM games to include combinatorial action spaces for learner, which might be exponentially large, and infinite action spaces for the adversary. Neither of these situations can be handled by generic algorithms for finite PM games. Specifically, the model considered an action space $\mathcal{X}$ for the learner, that has a small subset of actions defining a *global observable set* (see Assumption 2 in Section 2). The adversary's action space is a continuous, bounded vector space with the adversary sampling moves from a fixed distribution over the vector space. The reward function considered is a general non-linear function of learner's and adversary's actions, with some restrictions (see Assumptions 1 & 3 in Section 2). The model incorporated a linear feedback mechanism where the feedback received is a linear transformation of adversary's move. Inspired by the classic confidence bound algorithms for MABs, such as UCB [6], the authors proposed a Global Confidence Bound (GCB) algorithm that enjoyed two types of regret bound. The first one was a distribution independent $O(T^{2/3}\log T)$ regret bound and the second one was a distribution dependent $O(\log T)$ regret bound. A distribution dependent regret bound involves factors specific to the adversary's fixed distribution, while distribution independent means the regret bound holds over all possible distributions in a broad class of distributions. Both bounds also had a logarithmic dependence on $|\mathcal{X}|$. The algorithm combined online estimation with two offline computational oracles. The first oracle finds the action(s) achieving maximum value of reward function over $\mathcal{X}$, for a particular adversary action (argmax oracle), and the second oracle finds the action(s) achieving second maximum value of reward function over $\mathcal{X}$, for a particular adversary action (arg-secondmax oracle). Moreover, the distribution dependent regret bound requires existence of a *unique* optimal learner action. The inspiration for the CPM model came from various applications like crowdsourcing and matching problems like matching products with customers.

Our Contributions. We adopt the CPM model proposed earlier [1]. However, instead of using upper confidence bound techniques, our work is motivated by another classic technique developed for MABs, namely that of forced exploration. This technique was already used in the classic paper of Robbins [9] and has also been called "forcing with certainty equivalence" in the control theory literature [10]. We develop a Phased Exploration with Greedy Exploitation (PEGE) algorithmic framework (Section 3) borrowing the PEGE terminology from work on linearly parameterized bandits [11]. When the framework is instantiated with different parameters, it achieves $O(T^{2/3}\sqrt{\log T})$ distribution independent and $O(\log^2 T)$ distribution dependent regret. Significantly, the framework combines online estimation with only the argmax oracle from GCB, which is a practical advantage over requiring an additional arg-secondmax oracle. Moreover, the distribution dependent regret does not require existence of unique optimal action. Uniqueness of optimal action can be an unreasonable assumption, especially in the presence of a combinatorial action space. Our second contribution is another algorithm PEGE2 (Section 4) that combines a PEGE algorithm with Gap Estimation, to achieve a distribution dependent $O(\log T)$ regret bound, thus matching the GCB regret guarantee in terms of T and gap. Here, gap refers to the difference between expected reward of optimal and second optimal learner's actions. However, PEGE2 does require access to both the oracles and existence of unique optimal action. On the other hand, like GCB, PEGE2 has the property that its regret is never larger than $O(T^{2/3}\log T)$ even when there is no unique optimal action. Another crucial fact

is that all our regret bounds are independent of $|\mathcal{X}|$, only depending on the size of the small *global observable set*. Thus, though we have adopted the CPM model [1], our regret bounds are meaningful for countably infinite or even continuous learner's action space, whereas GCB regret bound has an explicit logarithmic dependence on $|\mathcal{X}|$. We provide a detailed comparison of our work with the GCB algorithm in Section 5. Finally, we discuss how our algorithms can be efficiently applied in the CPM problem of online ranking with feedback restricted to top ranked items (Section 6), a setting already considered [12] but analyzed in a non-stochastic setting.

2 Preliminaries and Assumptions

The online game is played between a learner and an adversary, over discrete rounds indexed by $t = 1, 2, \ldots$. The learner's action set is denoted as $\mathcal{X}$ which can be exponentially large. The adversary's action set is the infinite set $[0,1]^n$. The adversary fixes a distribution p on $[0,1]^n$ before start of the game (adversary's strategy), with p unknown to the learner. At each round of the game, adversary samples $\theta(t) \in [0,1]^n$ according to p, with $\mathbb{E}_{\theta(t)\sim p}[\theta(t)] = \theta^*_p$. The learner chooses $x(t) \in \mathcal{X}$ and gets reward $r(x(t), \theta(t))$. However, the learner might not get to know either $\theta(t)$ (as in a full information game) or $r(x(t), \theta(t))$ (as in a bandit game). In fact, the learner receives, as feedback, a linear transformation of $\theta(t)$.That is, every action $x \in \mathcal{X}$ has an associated transformation matrix $M_x \in \mathbb{R}^{m_x \times n}$. On playing action $x(t)$, the learner receives a feedback $M_{x(t)} \cdot \theta(t) \in \mathbb{R}^{m_x}$. Note that the game with the defined feedback mechanism subsumes full information and bandit games. $M_x = \mathbb{I}^{n\times n}$, $\forall x$ makes it a full information game since $M_x \cdot \theta = \theta$. If $r(x, \theta) = x \cdot \theta$, then $M_x = x \in \mathbb{R}^n$ makes it a bandit game. The dimension n, action space $\mathcal{X}$, reward function $r(\cdot, \cdot)$ and transformation matrices M_x, $\forall x \in \mathcal{X}$ are known to the learner. The goal of the learner is to minimize the expected regret, which, for a given time horizon T, is:

$$R(T) = T \cdot \max_{x\in\mathcal{X}} \mathbb{E}[r(x, \theta)] - \sum_{t=1}^{T} \mathbb{E}[r(x(t), \theta(t))] \tag{1}$$

where the expectation in the first term is taken over θ, w.r.t. distribution p, and the second expectation is taken over θ and possible randomness in the learner's algorithm.

For distribution dependent regret bounds, we define gaps in expected rewards: Let $x^* \in S(x) = \text{argmax}_{x\in X}\, \bar{r}(x, \theta^*_p)$. Then $\Delta_x = \bar{r}(x^*, \theta^*_p) - \bar{r}(x, \theta^*_p)$, $\Delta_{max} = \max\{\Delta_x : x \in \mathcal{X}\}$ and $\Delta = \min\{\Delta_x : x \in \mathcal{X}, \Delta_x > 0\}$.

Assumption 1. (Restriction on Reward Function) The first assumption is that $\mathbb{E}_{\theta\sim p}[r(x, \theta)] = \bar{r}(x, \theta^*_p)$, for some function $\bar{r}(\cdot, \cdot)$. That is, the expected reward is a function of x and θ^*_p, which is always satisfied if $r(x, \theta)$ is a linear function of θ, or if distribution p happens to be any distribution with support $[0,1]^n$ and fully parameterized by its mean θ^*_p. With this assumption, the expected regret becomes:

$$R(T) = T \cdot \bar{r}(x^*, \theta^*_p) - \sum_{t=1}^{T} \mathbb{E}[\bar{r}(x(t), \theta^*_p)]. \tag{2}$$

Assumption 2. (Existence of Global Observable Set) The second assumption is on the existence of a *global observable set*, which is a subset of learner's action set and is required for estimating an adversary's move θ. The *global observable set* is defined as follows: for a set of actions $\sigma = \{x_1, x_2, \ldots, x_{|\sigma|}\} \subseteq \mathcal{X}$, let their transformation matrices be stacked in a top down fashion to obtain a $\mathbb{R}^{\sum_{i=1}^{|\sigma|} m_{x_i} \times n}$ dimensional matrix M_σ. σ is said to be a global observable set if M_σ has full column rank, i.e., rank(M_σ) = n. Then, the Moore-Penrose pseudoinverse M^+_σ satisfies $M^+_\sigma M_\sigma = \mathbb{I}^{n\times n}$. Without the assumption on the existence of global observable set, it might be the case that even if the learner plays all actions in $\mathcal{X}$ on same θ, the learner might not be able to recover θ (as $M^+_\sigma M_\sigma = \mathcal{I}^{n\times n}$ will not hold without full rank assumption). In that case, learner might not be able to distinguish between $\theta^*_{p_1}$ and $\theta^*_{p_2}$, corresponding to two different adversary's strategies. Then, with non-zero probability, the learner can suffer $\Omega(T)$ regret and no learner strategy can guarantee a sub-linear in T regret (the intuition forms the base of the *global observability condition* in [2]). Note that the size of the global observable set is small, i.e., $|\sigma| \leq n$. A global observable set can be found by including an action x in σ if it strictly increases the rank of M_σ, till the rank reaches n. There can, of course, be more than one global observable set.

Assumption 3. (Lipschitz Continuity of Expected Reward Function) The third assumption is on the Lipschitz continuity of expected reward function in its second argument. More precisely, it is assumed that $\exists\ R > 0$ such that $\forall\ x \in \mathcal{X}$, for any θ_1 and θ_2, $|\bar{r}(x, \theta_1) - \bar{r}(x, \theta_2)| \leq R\|\theta_1 - \theta_2\|_2$. This assumption is reasonable since otherwise, a small error in estimation of mean reward vector θ_p^* can introduce a large change in expected reward, leading to difficulty in controlling regret over time. The Lipschitz condition holds trivially for expected reward functions which are linear in second argument. The continuity assumption, along with the fact that adversary's moves are in $[0, 1]^n$, implies boundedness of expected reward for any learner's action and any adversary's action. We denote $R_{max} = \max_{x \in \mathcal{X}, \theta \in [0,1]^n} \bar{r}(x, \theta)$.

The three assumptions above will be made throughout. However, the fourth assumption will only be made in a subset of our results.

Assumption 4. (Unique Optimal Action) The optimal action $x^* = \text{argmax}_{x \in \mathcal{X}}\, \bar{r}(x, \theta_p^*)$ is unique. Denote a second best action (which may not be unique) by $x_-^* = \text{argmax}_{x \in \mathcal{X}, x \neq x^*}\, \bar{r}(x, \theta_p^*)$. Note that $\Delta = \bar{r}(x^*, \theta_p^*) - \bar{r}(x_-^*, \theta_p^*)$.

3 Phased Exploration with Greedy Exploitation

Algorithm 1 (PEGE) uses the classic idea of doing exploration in phases that are successively further apart from each other. In between exploration phases, we select action greedily by completely trusting the current estimates. The constant β controls how much we explore in a given phase and the constant α along with the function $C(\cdot)$ determines how much we exploit. This idea is classic in the bandit literature [9–11] but has not been applied to the CPM framework to the best of our knowledge.

Algorithm 1 The PEGE Algorithmic Framework

Inputs: α, β and function $C(\cdot)$ (to determine amount of exploration/exploitation in each phase).

For $b = 1, 2, \ldots,$
  Exploration
    For $i = 1$ to $|\sigma|$
      For $j = 1$ to b^β
        Let $t_{j,i} = t$ and $\theta(t_{j,i}, b) = \theta(t)$ where t is current time point
        Play $x_i \in \sigma$ and get feedback $M_{x_i} \cdot \theta(t_{j,i}, b) \in \mathbb{R}^{m_{x_i}}$.
      End For
    End For
  Estimation
  $\tilde{\theta}_{j,i} = M_\sigma^+(M_{x_1} \cdot \theta(t_{j,1}, i), \ldots, M_{x_{|\sigma|}} \cdot \theta(t_{j,|\sigma|}, i)) \in \mathbb{R}^n$.
  $\hat{\theta}(b) = \dfrac{\sum_{i=1}^{b} \sum_{j=1}^{i^\beta} \tilde{\theta}_{j,i}}{\sum_{j=1}^{b} j^\beta} \in \mathbb{R}^n$.
  $x(b) \in \text{argmax}_{x \in \mathcal{X}}\, \bar{r}(x, \hat{\theta}(b))$.
  Exploitation
    For $i = 1$ to $\exp(C(b^\alpha))$
      Play $x(b)$.
    End For
End For

It is easy to see that the estimators in Algorithm 1 have the following properties: $\mathbb{E}_p[\tilde{\theta}_{j,i}] = M_\sigma^+(M_{x_1} \cdot \theta_p^*, \ldots, M_{x_{|\sigma|}} \cdot \theta_p^*) = M_\sigma^+ M_\sigma \cdot \theta_p^* = \theta_p^*$ and hence $\mathbb{E}_p[\hat{\theta}] = \theta_p^*$. Using the fact that

$M_\sigma^+ = (M_\sigma^\top M_\sigma)^{-1} M_\sigma^\top$, we also have the following bound on estimation error of θ_p^*:

$$\begin{aligned}
&\|\tilde{\theta}_{j,i} - \theta_p^*\|_2 \leq \|M_\sigma^+ (M_{x_1} \cdot \theta(t_{j,1}, i), \ldots, M_{x_{|\sigma|}} \cdot \theta(t_{j,|\sigma|}, i)) - M_\sigma^+ M_\sigma \theta_p^*\|_2 \\
&= \|(M_\sigma^\top M_\sigma)^{-1} \sum_{k=1}^{|\sigma|} M_{x_k}^\top M_{x_k} \cdot (\theta(t_{j,k}, i) - \theta_p^*)\|_2 \leq \sqrt{n} \sum_{k=1}^{|\sigma|} \|(M_\sigma^\top M_\sigma)^{-1} M_{x_k}^\top M_{x_k}\|_2 =: \beta_\sigma
\end{aligned} \tag{3}$$

where the constant β_σ defined above depends only on the structure of the linear transformation matrices of the global observer set and not on adversary strategy p.

Our first result is about the regret of Algorithm 1 when within phase number b, the exploration part spends $|\sigma|$ rounds (constant w.r.t. b) and the exploitation part grows polynomially with b.

Theorem 1. (Distribution Independent Regret) *When Algorithm 1 is initialized with the parameters $C(a) = \log a$, $\alpha = 1/2$ and $\beta = 0$, and the online game is played over T rounds, we get the following bound on expected regret:*

$$R(T) \leq R_{max} |\sigma| T^{2/3} + 2R\beta_\sigma T^{2/3} \sqrt{\log 2e^2 + 2 \log T} + R_{max} \tag{4}$$

where β_σ is the constant as defined in Eq. 3.

Our next result is about the regret of Algorithm 1 when within phase number b, the exploration part spends $|\sigma| \cdot b$ rounds (linearly increasing with b) and the exploitation part grows exponentially with b.

Theorem 2. (Distribution Dependent Regret) *When Algorithm 1 is initialized with the parameters $C(a) = h \cdot a$, for a tuning parameter $h > 0$, $\alpha = 1$ and $\beta = 1$, and the online game is played over T rounds, we get the following bound on expected regret:*

$$R(T) \leq \sum_{x \in \sigma} \Delta_x \left(\frac{\log T}{h} \right)^2 + \frac{4\sqrt{2\pi} e^2 R \Delta_{max} \beta_\sigma}{\Delta} e^{\frac{h^2 (2R^2 \beta_\sigma^2)}{\Delta^2}}. \tag{5}$$

Such an explicit bound for a PEGE algorithm that is polylogarithmic in T and explicitly states the multiplicative and additive constants involved in not known, to the best of our knowledge, even in the bandit literature (e.g., earlier bounds [10] are asymptotic) whereas here we prove it in the CPM setting. Note that the additive constant above, though finite, blows up exponentially fast as $\Delta \to 0$ for a fixed h. It is well behaved however, if the tuning parameter h is on the same scale as Δ. This line of thought motivates us to estimate the gap to within constant factors and then feed that estimate into a PEGE algorithm. This is what we will do in the next section.

4 Combining Gap Estimation with PEGE

Algorithm 2 tries to estimate the gap Δ to within a constant multiplicative factor. However, if there is no unique optimal action or when the true gap is small, gap estimation can take a very large amount of time. To prevent that from happening, the algorithm also takes in a threshold T_0 as input and definitely stops if the threshold is reached. The result below assures us that, with high probability, the algorithm behaves as expected. That is, if there is a unique optimal action and the gap is large enough to be estimated with a given confidence before the threshold T_0 kicks in, it will output an estimate $\hat{\Delta}$ in the range $[0.5\Delta, 1.5\Delta]$. On the other hand, if there is no unique optimal action, it does not generate an estimate of Δ and instead runs out of the exploration budget T_0.

Theorem 3. (Gap Estimation within Constant Factors) *Let $T_0 \geq 1$ and $\delta \in (0, 1)$ and define $T_1(\delta) = \frac{256 R^2 \beta_\sigma^2}{\Delta^2} \log \frac{512 e^2 R^2 \beta_\sigma^2}{\Delta^2 \delta}$, $T_2(\delta) = \frac{16 R^2 \beta_\sigma^2}{\Delta^2} \log \frac{4e^2}{\delta}$. Consider Algorithm 2 run with*

$$w(b) = \sqrt{\frac{R^2 \beta_\sigma^2 \log(\frac{4e^2 b^2}{\delta})}{b}}. \tag{6}$$

Then, the following 3 claims hold.

1. *Suppose Assumption 4 holds and $T_1(\delta) < T_0$. Then with probability at least $1 - \delta$, Algorithm 2 stops in $T_1(\delta)$ episodes and outputs an estimate $\hat{\Delta}$ that satisfies $\frac{1}{2}\Delta \leq \hat{\Delta} \leq \frac{3}{2}\Delta$.*

Algorithm 2 Algorithm for Gap Estimation

Inputs: T_0 (exploration threshold) and δ (confidence parameter)

For b = 1, 2, ...,
  Exploration
    For i = 1 to $|\sigma|$
      (Denote) $t_i = t$ and $\theta(t_i, b) = \theta(t)$ (t is current time point).
      Play $x_i \in \sigma$ and get feedback $M_{x_i} \cdot \theta(t_i, b) \in \mathbb{R}^{m_{x_i}}$.
    End For
  Estimation
  $\tilde{\theta}_b = M_\sigma^+(M_{x_1} \cdot \theta(t_1, b), \ldots, M_{x_{|\sigma|}} \cdot \theta(t_{|\sigma|}, b)) \in \mathbb{R}^n$.
  $\hat{\theta}(b) = \dfrac{\sum_{i=1}^{b} \tilde{\theta}_i}{b} \in \mathbb{R}^n$.

  Stopping Rule ($w(b)$ is defined as in Eq. (6))
  If $\operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \hat{\theta}(b))$ is unique:
    $\hat{x}(b) = \operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \hat{\theta}(b))$
    $\hat{x}_-(b) = \operatorname{argmax}_{x \in \mathcal{X}, x \neq \hat{x}(b)} \bar{r}(x, \hat{\theta}(b))$ (need not be unique)
    If $\bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) > 6w(b)$:
      STOP and output $\hat{\Delta} = \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b))$
    End If
  End If
  If $b > T_0$:
    STOP and output "threshold exceeded"
  End If
End For

2. *Suppose Assumption 4 holds and $T_0 \leq T_1(\delta)$. Then with probability at least $1 - \delta$, the algorithm either outputs "threshold exceeded" or outputs an estimate $\hat{\Delta}$ that satisfies $\frac{1}{2}\Delta \leq \hat{\Delta} \leq \frac{3}{2}\Delta$. Furthermore, if it outputs $\hat{\Delta}$, it must be the case that the algorithm stopped at an episode b such that $T_2(\delta) < b < T_0$.*

3. *Suppose Assumption 4 fails. Then, with probability at least $1 - \delta$, Algorithm 2 stops in T_0 episodes and outputs "threshold exceeded".*

Equipped with Theorem 3, we are now ready to combine Algorithm 2 with Algorithm 1 to give Algorithm 3. Algorithm 3 first calls Algorithm 2. If Algorithm 2 outputs an estimate $\hat{\Delta}$ it is fed into Algorithm 1. If the threshold T_0 is exceeded, then the remaining time is spent in pure exploitation. Note that by choosing T_0 to be of order $T^{2/3}$ we can guarantee a worst case regret of the same order even when unique optimality assumption fails. For PM games that are globally observable but not locally observable, such a distribution independent $O(T^{2/3})$ bound is known to be optimal [4].

Theorem 4. (Regret Bound for PEGE2) *Consider Algorithm 3 run with knowledge of the number T of rounds. Consider the distribution independent bound*

$$B_1(T) = 2(2R\beta_\sigma|\sigma|^2 R_{max}^2 T)^{2/3}\sqrt{\log(4e^2T^3)} + R_{max},$$

and the distribution dependent bound

$$B_2(T) = \frac{256R^2\beta_\sigma^2}{\Delta^2}\log\frac{512e^2R^2\beta_\sigma^2 T}{\Delta^2}R_{max}|\sigma| + \sum_{x\in\sigma}\Delta_x\frac{36R^2\beta_\sigma^2\log T}{\Delta^2} + \frac{8e^2R^2\beta_\sigma^2}{\Delta^2} + R_{max}.$$

If Assumption 4 fails, then the expected regret of Algorithm 3 is bounded as $R(T) \leq B_1(T)$. If Assumption 4 holds, then the expected regret of Algorithm 3 is bounded as

$$R(T) \leq \begin{cases} B_2(T) & \text{if } T_1(\delta) < T_0 \\ O(T^{2/3}\log T) & \text{if } T_0 \leq T_1(\delta) \end{cases}, \tag{7}$$

where $T_1(\delta)$ is as defined in Theorem 3 and δ, T_0 are as defined in Algorithm 3.

In the above theorem, note that $T_1(\delta)$ scales as $\Theta(\frac{1}{\Delta^2}\log\frac{T}{\Delta^2})$ and T_0 as $\Theta(T^{2/3})$. Thus, the two cases in Eq. (7) correspond to large gap and small gap situations respectively.

Algorithm 3 Algorithm Combining PEGE with Gap Estimation (PEGE2)

1: Input: T (total number of rounds)

2: Call Algorithm 2 with inputs $T_0 = \left(\frac{2R\beta_\sigma T}{|\sigma|R_{max}}\right)^{2/3}$ and $\delta = 1/T$
3: If Algorithm 2 returns "threshold exceeded":
4: Let $\hat{\theta}(T_0)$ be the latest estimate of θ_p^* maintained by Algorithm 2
5: Play $\hat{x}(T_0) = \text{argmax}_{x\in\mathcal{X}}\,\bar{r}(x,\hat{\theta})$ for the remaining $T - T_0|\sigma|$ rounds
6: Else:
7: Let $\hat{\Delta}$ be the gap estimate produced by Algorithm 2
8: For all remaining time steps, run Algorithm 1 with parameters $C(a) = ha$ with $h = \frac{\hat{\Delta}^2}{9R^2\beta_\sigma^2}$, $\alpha = 1$, $\beta = 0$
9: End If

5 Comparison with GCB Algorithm

We provide a detailed comparison of our results with those obtained for GCB [1]. (a) While we use the same CPM model, our work is inspired by the forced exploration technique while GCB is inspired by the confidence bound technique, both of which are classic in the bandit literature. (b) One instantiation of our PEGE framework gives an $O(T^{2/3}\sqrt{\log T})$ distribution independent regret bound (Theorem 1), which does not require call to arg-secondmax oracle. This is of substantial practical advantage over GCB since even for linear optimization problems over polyhedra, standard routines usually do not have option of computing action(s) that achieve second maximum value for the objective function. (c) Another instantiation of the PEGE framework gives an $O(\log^2 T)$ distribution dependent regret bound (Theorem 2), which neither requires call to arg-secondmax oracle nor the assumption of existence of unique optimal action for learner. This is once again important, since the assumption of existence of unique optimal action might be impractical, especially for exponentially large action space. However, the caveat is that improper setting of the tuning parameter h in Theorem 2 can lead to a large additive component in the regret. (d) A crucial point, which we had highlighted in the beginning, is that the regret bounds achieved by PEGE and PEGE2 do not have dependence on size of learner's action space, i.e., $|\mathcal{X}|$. The dependence is only on the size of global observable set σ, which is guaranteed to be not more than dimension of adversary's action space. Thus, though we have adopted the CPM model [1], our algorithms achieve meaningful regret bounds for countably infinite or even continuous learner's action space. In contrast, the GCB regret bounds have explicit, logarithmic dependence on size of learner's action space. Thus, their results cannot be extended to problems with infinite learner's action space (see Section 6 for an example), and are restricted to large, but *finite* action spaces. (e) The PEGE2 algorithm is a true analogue of the GCB algorithm, matching the regret bounds of GCB in terms of T and gap Δ with the advantage that it has no dependence on $|\mathcal{X}|$. The disadvantage, however, is that PEGE2 requires knowledge of time horizon T, while GCB is an anytime algorithm.

6 Application to Online Ranking

A recent paper studied an interesting problem of online ranking with feedback restricted to top ranked items [12]. The problem was studied in a non-stochastic setting, i.e., it was assumed that an oblivious adversary generates reward vectors. Moreover, the learner's action space was exponentially large in number of items to be ranked. The paper made the connection of the problem setting to PM games (but not combinatorial PM games) and proposed an efficient algorithm for the specific problem at hand. However, a careful reading of the paper shows that their algorithmic techniques can handle the CPM model we have discussed so far, but in the *non-stochastic* setting. The reward function is linear in both learner's and adversary's moves, adversary's move is restricted to a finite space of vectors and feedback is a linear transformation of adversary's move. In this section, we give

a brief description of the problem setting and show how our algorithms can be used to efficiently solve the problem of online ranking with feedback on top ranked items in the *stochastic* setting. We also give an example of how the ranking problem setting can be somewhat naturally extended to one which has continuous action space for learner, instead of large but finite action space.

The paper considered an online ranking problem, where a learner repeatedly re-ranks a set of n, fixed items, to satisfy diverse users' preferences, who visit the system sequentially. Each learner action x is a permutation of the n items. Each user has like/dislike preference for each item, varying between users, with each user's preferences encoded as n length binary relevance vectors θ. Once the ranked list of items is presented to the user, the user scans through the items, but gives relevance feedback only on top ranked item. However, the performance of the learner is judged based on full ranked list and unrevealed, full relevance vector. Thus, we have a PM game, where neither adversary generated relevance vector nor reward is revealed to learner. The paper showed how a number of practical ranking measures, like Discounted Cumulative Gain (DCG), can be expressed as a linear function, i.e., $r(x, \theta) = f(x) \cdot \theta$. The practical motivation of the work was based on learning a ranking strategy to satisfy diverse user preferences, but with limited feedback received due to user burden constraints and privacy concerns.

Online Ranking with Feedback at Top as a Stochastic CPM Game. We show how our algorithms can be applied in online ranking with feedback for top ranked items by showing how it is a specific instance of the CPM model and how our key assumptions are satisfied. The learner's action space is the finite but exponentially large space of $\mathcal{X} = n!$ permutations. Adversary's move is an n dimensional relevance vector, and thus, is restricted to $\{0, 1\}^n$ (finite space of size 2^n) contained in $[0, 1]^n$. In the stochastic setting, we can assume that adversary samples $\theta \in \{0, 1\}^n$ from a fixed distribution on the space. Since the feedback on playing a permutation is the relevance of top ranked item, each move x has an associated transformation matrix $M_x \in \{0, 1\}^n$, with 1 in the place of the item which is ranked at the top by x and 0 everywhere else. Thus, $M_x \cdot \theta$ gives the relevance of item ranked at the top by x. The global observable set σ is the set of any n actions, where each action, in turn, puts a distinct item on top. Hence, M_σ is the $n \times n$ dimensional permutation matrix. Assumption 1 is satisfied because the reward function is linear in θ and $\bar{r}(x, \theta^*_p) = f(x) \cdot \theta^*_p$, where $\mathbb{E}_p[\theta] = \theta^*_p \in [0, 1]^n$. Assumption 2 is satisfied since there will always be a global observable set of size n and can be found easily. In fact, there will be multiple global observable sets, with the freedom to choose any one of them. Assumption 3 is satisfied due to the expected reward function being linear in second argument. The Lipschitz constant is $\max_{x \in \mathcal{X}} \|f(x)\|_2$, which is always less than some small polynomial factor of n, depending on specific $f(\cdot)$. The value of β_σ can be easily seen to be $n^{3/2}$. The argmax oracle returns the permutation which simply sorts items according to their corresponding θ values. The arg-secondmax oracle is more complicated, though feasible. It requires first sorting the items according to θ and then compare each pair of consecutive items to see where least drop in reward value occurs.

Likely Failure of Unique Optimal Action Assumption. Assumption 4 is unlikely to hold in this problem setting. The mean relevance vector θ^*_p effectively reflects the average preference of all users for each of the n items. It is very likely that at least a few items will not be liked by anyone and which will ultimately be always ranked at the bottom. Equally possible is that two items will have same user preference on average, and can be exchanged without hurting the optimal ranking. Thus, existence of an unique optimal ranking, which indicates that each item will have different average user preference than every other item, is unlikely. Thus, PEGE algorithm can still be applied to get poly-logarithmic regret (Theorem 2), but GCB will only achieve $O(T^{2/3} \log T)$ regret.

A PM Game with Infinite Learner Action Space. We give a simple modification of the ranking problem above to show how the learner can have continuous action space. The learner now ranks the items by producing an n dimensional score vector $x \in [0, 1]^n$ and sorting items according to their scores. Thus the learner's action space is now an uncountably infinite continuous space. As before, the user gets to see the ranked list and gives relevance feedback on top ranked item. The learner's performance will now be judged by a continuous loss function, instead of a discrete-valued ranking measure, since its moves are in a continuous space. Consider the simplest loss, viz., the squared "loss" $r(x, \theta) = -\|x - \theta\|_2^2$ (note -ve sign to keep reward interpetation). It can be easily seen that $\bar{r}(x, \theta^*_p) = E_{\theta \sim p}[r(x, \theta)] = -\|x\|_2^2 + 2x \cdot \theta^*_p - \mathbf{1} \cdot \theta^*_p$, if the relevance vectors θ are in $\{0, 1\}^n$. Thus, the Lipschitz condition is satisfied. The global observable set is still of size n, with the n actions being any n score vectors, whose sorted orders place each of the n items, in turn, on top. β_σ remains same as before, with $\mathrm{argmax}_x \, \mathbb{E}_{\theta \sim p} r(x, \theta) = \mathbb{E}_{\theta \sim p}[\theta] = \theta^*_p$.

References

[1] Tian Lin, Bruno Abrahao, Robert Kleinberg, John Lui, and Wei Chen. Combinatorial partial monitoring game with linear feedback and its applications. In *Proceedings of the 31th International Conference on Machine Learning*, pages 901–909. ACM, 2014.

[2] Antonio Piccolboni and Christian Schindelhauer. Discrete prediction games with arbitrary feedback and loss. In *COLT*, pages 208–223. Springer, 2001.

[3] Nicolo Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, pages 562–580, 2006.

[4] Gabor Bartok et al. Partial monitoring–classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997, 2014.

[5] Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. In *Advances in Neural Information Processing Systems*, pages 1783–1791, 2015.

[6] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.

[7] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *Proceedings of the 30th International Conference on Machine Learning*, pages 151–159, 2013.

[8] Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. Tight regret bounds for stochastic combinatorial semi-bandits. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pages 535–543, 2015.

[9] Herbert Robbins. Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pages 169–177. Springer, 1985.

[10] Rajeev Agrawal and Demosthenis Teneketzis. Certainty equivalence control with forcing: revisited. *Systems & control letters*, 13(5):405–412, 1989.

[11] Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.

[12] Sougata Chaudhuri and Ambuj Tewari. Online ranking with top-1 feedback. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, pages 129–137. ACM, 2015.

[13] Thomas P Hayes. A large-deviation inequality for vector-valued martingales. *Combinatorics, Probability and Computing*, 2005.

7 Appendix

We first state the large deviation inequality for vector-valued martingales, which is the generalization of Azuma-Hoeffding inequality for scalar valued martingales.

Theorem 1.8 of [13]: Let $X_0, X_1, \ldots, X_m$ be a weak martingale sequence taking values in euclidean space $\mathbb{R}^d$, with $\mathbb{E}[X_i|X_{i-1}] = X_{i-1}$. Let $X_0 = 0$ and $\|X_i - X_{i-1}\|_2 \leq 1$, for $i = 1, 2, \ldots, m$. Then, for every $\epsilon > 0$,

$$\Pr[\|X_m\|_2 \geq \epsilon] < 2e^2 e^{\frac{-\epsilon^2}{m}} \tag{8}$$

We use the concentration inequality to get a uniform confidence bound, over the space of learner's action, on the deviation of estimated reward from true reward, after each estimate of mean reward vector is produced.

Lemma 5. *At the end of exploration phase within phase b, $b = 1, 2, \ldots$, of Algorithm PEGE, the estimator of reward vector θ_p^* is $\hat{\theta}(b) = \dfrac{\sum_{i=1}^{b}\sum_{j=1}^{i^\beta}\tilde{\theta}_{j,i}}{\sum_{j=1}^{b} j^\beta}$. Then, $\forall\ \eta > 0$,*

$$\Pr[\forall\ x \in \mathcal{X} : |\bar{r}(x, \hat{\theta}(b)) - \bar{r}(x, \theta_p^*)| \leq \eta] \geq 1 - 2e^2 e^{\frac{-(\sum_{i=1}^{b^\beta} i^\beta)\eta^2}{R^2\beta_\sigma^2}} \tag{9}$$

where β_σ is the constant as defined in Eq. 3 and R is the Lipschitz constant defined in Assumption 3.

Proof. Let $\{X_{i,j}\}_{\substack{j=1,\ldots,i^\beta \\ i=1,\ldots,b}}$ be a sequence of random vectors, defined as follows:

$$X_{i,j} = \frac{\sum_{i'=1}^{i-1}\sum_{j'=1}^{i'^\beta}\theta_p^* + \sum_{k=1}^{j}\theta_p^* - (\sum_{i'=1}^{i-1}\sum_{j'=1}^{i'^\beta}\tilde{\theta}_{j',i'} + \sum_{k=1}^{j}\tilde{\theta}_{k,i})}{\sum_{i''=1}^{b}\sum_{j''=1}^{(i'')^\beta}\beta_\sigma} \tag{10}$$

It can be checked that the ℓ_2 norm of the difference between any two consecutive random vectors is bounded by a constant. That is, $\|X_{i,j} - X_{i,j-1}\|_2 = \dfrac{\|\theta_p^* - \tilde{\theta}_{j,i}\|_2}{\sum_{i''=1}^{b}\sum_{j''=1}^{(i'')^\beta}\beta_\sigma} \leq \dfrac{1}{\sum_{i''=1}^{b}(i'')^\beta}$ and $\|X_{i+1,1} - X_{i,i^\beta}\|_2 \dfrac{\|\theta_p^* - \tilde{\theta}_{1,i+1}\|_2}{\sum_{i''=1}^{b}\sum_{j''=1}^{(i'')^\beta}\beta_\sigma} \leq \dfrac{1}{\sum_{i''=1}^{b}(i'')^\beta}$.

Also, $\tilde{\theta}_{j,i}$ is independent of all estimators formed before $\tilde{\theta}_{j,i}$ in Algorithm PEGE. Thus,

$$\begin{aligned}\mathbb{E}[X_{i,j} - X_{i,j-1}|X_{i,j-1}] &= \mathbb{E}\left[\frac{\theta_p^* - \tilde{\theta}_{j,i}}{\sum_{i''=1}^{b}\sum_{j''=1}^{(i'')^\beta}\beta_\sigma}|X_{i,j-1}\right] \\ &= \mathbb{E}\left[\frac{\theta_p^* - \tilde{\theta}_{j,i}}{\sum_{i''=1}^{b}\sum_{j''=1}^{(i'')^\beta}\beta_\sigma}\right] = 0\end{aligned} \tag{11}$$

Thus, $\{X_{i,j}\}_{\substack{j=1,\ldots,i^\beta \\ i=1,\ldots,b}}$ satisfy the criteria of weak martingale sequence and hence, by the large deviation inequality of vector valued martingales, we have, $\forall\ \epsilon > 0$, $\Pr[\|X_{b,b^\beta}\|_2 \geq \epsilon] < 2e^2 e^{\left(\frac{-\epsilon^2}{\frac{\sum_{i=1}^{b}\sum_{j=1}^{i^\beta}1}{(\sum_{i=1}^{b} i^\beta)^2}}\right)} = 2e^2 e^{-\epsilon^2\sum_{i=1}^{b} i^\beta}$.

Now, it can be clearly seen that $\|\theta_p^* - \hat{\theta}(b)\|_2 = \beta_\sigma\|X_{b,b^\beta}\|_2$ and let $\eta = \beta_\sigma\epsilon$. Then, $\forall\ \eta > 0$, we get $\Pr[\|\theta_p^* - \hat{\theta}(b)\|_2 \geq \eta] \leq 2e^2 e^{\frac{-(\sum_{i=1}^{b^\beta} i^\beta)\eta^2}{\beta_\sigma^2}}$.

Using the Lipschitz property of expected reward function (Assumption 3), we have

$$\begin{aligned}\Pr(\exists\, x \in \mathcal{X} : |\bar{r}(x, \hat{\theta}(b)) - \bar{r}(x, \theta_p^*)| \geq \eta) &\leq \Pr(R \cdot \|\theta_p^* - \hat{\theta}(b)\|_2 \geq \eta) \\ &\leq 2e^2 e^{\frac{-(\sum_{i=1}^{b^\beta} i^\beta)\eta^2}{R^2\beta_\sigma^2}}\end{aligned} \tag{12}$$

Taking complement of the event completes the proof. □

7.1 Proof of Results in Section 3

7.1.1 Proof of Theorem 1

We first restate the theorem.

Distribution Independent Regret: When Algorithm PEGE is initialized with the parameters $C(a) = \log a$, $\alpha = 1/2$ and $\beta = 0$, and the online game is played over T rounds, we get the following bound on expected regret:

$$R(T) \leq R_{max}|\sigma|T^{2/3} + 2R\beta_\sigma T^{2/3}\sqrt{\log 2e^2 + 2\log T} + R_{max} \tag{13}$$

where β_σ is the constant as defined in Eq. 3.

Proof. Let Algorithm PEGE run for K phases, with parameters initialized as $C(a) = \log a$, $\alpha = 1/2$ and $\beta = 0$.

Exploration regret: During every exploration phase, the expected regret is bounded by $|\sigma|R_{max}$, where R_{max} is as given in Assumption 3. Thus, total expected regret due to exploration is $K|\sigma|R_{max}$.

Exploitation regret: Let $x^* \in \operatorname{argmax}_{x\in\mathcal{X}} \bar{r}(x, \theta_p^*)$ and $x(b) \in \operatorname{argmax}_{x\in\mathcal{X}} \bar{r}(x, \hat{\theta}(b))$. During every exploitation round within phase b of Algorithm PEGE, the expected regret is $|\bar{r}(x(b), \theta_p^*) - \bar{r}(x^*, \theta_p^*)|$.

Now, from Lemma 5, with $\beta = 0$, the following holds w.p. $\geq 1 - \delta_b$,

$$\forall\, x,\ |\bar{r}(x, \theta_p^*) - \bar{r}(x, \hat{\theta}(b)| \leq \underbrace{\sqrt{\frac{R^2\beta_\sigma^2 \log(\frac{2e^2}{\delta_b})}{b}}}_{\eta_b} \tag{14}$$

Then, w.p. $\geq 1 - \delta_b$, the following event holds true: $|\bar{r}(x^*, \theta_p^*) - \bar{r}(x(b), \theta_p^*)| \leq 2\eta_b$, as explained:

$$\begin{aligned}\bar{r}(x^*, \theta_p^*) &\leq \bar{r}(x^*, \hat{\theta}(b)) + \eta_b \ \text{ from Eq. } 14 \\ &\leq \bar{r}(x(b), \hat{\theta}(b)) + \eta_b \\ &\leq \bar{r}(x(b), \theta_p^*) + 2\eta_b \ \text{ from Eq. } 14\end{aligned}$$

Thus, the event $|\bar{r}(x^*, \theta_p^*) - \bar{r}(x(b), \theta_p^*)| \leq 2\eta_b$ holds true w.p. $\geq 1 - \delta_b$, for every fixed phase b. Then, w.p. $\geq 1 - \sum_{i=1}^K \delta_i$, the following holds true:

$$\forall\, b, |\bar{r}(x^*, \theta_p^*) - \bar{r}(x(b), \theta_p^*)| \leq 2\eta_b$$

Note that the expected regret per round is always bounded by $R_{\max}$ (since expected reward is bounded by R_{max}) .

The number of rounds of exploitation in phase b is b^α. Hence, the total expected regret due to exploitation, over K phases is:

$$\sum_{i=1}^{K} \left(\underbrace{(1 - \sum_{j=1}^{K} \delta_j)\frac{2R\beta_\sigma\sqrt{\log(2e^2/\delta_i)}}{\sqrt{i}} + (\sum_{j=1}^{K}\delta_j)R_{\max}}_{\text{expected regret per exploitation round}} \right) i^\alpha$$

Taking $\delta_1 = \delta_2 = \ldots = \delta_K = \delta$, and summing over exploration and exploitation regret over K phases, we get

$$R(T) \leq K|\sigma|R_{max} + \sum_{i=1}^{K} \left((1 - K\delta) \frac{2R\beta_\sigma \sqrt{\log(2e^2/\delta)}}{\sqrt{i}} + (K\delta) R_{\max} \right) i^\alpha \tag{15}$$

Using the inequality $\sum_{i=1}^{K} i^y \leq \int_0^K i^y dy \leq K^{y+1}$, we get expected regret:

$$R(T) \leq K|\sigma|R_{max} + (1 - K\delta) 2R\beta_\sigma \sqrt{\log(2e^2/\delta)} K^{\alpha+1/2} + K\delta R_{\max} K^{\alpha+1} \tag{16}$$

Now, we relate K to total time T as: $T = |\sigma|K + \sum_{i=1}^{K} i^\alpha \sim K^{\alpha+1}$, for large K.

Hence $K \sim T^{\frac{1}{1+\alpha}}$. Substituting value of K in Eq 16, and taking $\alpha = 1/2$ and $\delta = \frac{1}{KT}$ gives us the required bound on expected regret. □

Our next lemma shows that as the number of phases b grows in Algorithm PEGE, the probability of selecting a sub-optimal arm for greedy exploitation shrinks.

Lemma 6. *At the end of exploration phase within phase b, $b = 1, 2, \ldots$, the estimator constructed is $\hat{\theta}(b) = \dfrac{\sum_{i=1}^{b} \sum_{j=1}^{i^\beta} \tilde{\theta}_{j,i}}{\sum_{j=1}^{b} j^\beta}$. Then the following holds,*

$$\Pr(\operatorname*{argmax}_{x \in \mathcal{X}} \bar{r}(x, \hat{\theta}(b)) \not\subseteq \operatorname*{argmax}_{x \in \mathcal{X}} \bar{r}(x, \theta_p^*)) \leq 2e^2 e^{\frac{-(\sum_{i=1}^{b^\beta} i^\beta)\Delta^2}{4R^2\beta_\sigma^2}} \tag{17}$$

Proof. Let us assume $x' \in \operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \hat{\theta}(b))$ such that $x' \notin \operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \theta_p^*)$. Let $x^* \in \operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \theta_p^*)$. Then, by our assumption, $\bar{r}(x', \hat{\theta}(b)) \geq \bar{r}(x^*, \hat{\theta}(b))$. By definition of gap Δ, we also have $\bar{r}(x^*, \theta_p^*) - \bar{r}(x', \theta_p^*) \geq \Delta$. The two inequalities imply that at least one of the following two inequalities has to hold: either $|\bar{r}(x^*, \theta_p^*) - \bar{r}(x^*, \hat{\theta}(b))| \geq \frac{\Delta}{2}$ or $|\bar{r}(x', \hat{\theta}(b)) - \bar{r}(x', \theta_p^*)| \geq \frac{\Delta}{2}$.

Thus, $\operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \hat{\theta}(b)) \not\subseteq \operatorname{argmax}_{x \in \mathcal{X}} \bar{r}(x, \theta_p^*) \implies \exists\, x \in \mathcal{X} : |\bar{r}(x, \theta_p^*) - \bar{r}(x, \hat{\theta}(b))| \geq \frac{\Delta}{2}$. By using Lemma 5, and substituting $\eta = \frac{\Delta}{2}$, we get our result. □

7.1.2 Proof of Theorem 2

We restate the theorem before proving:

Distribution Dependent Regret: When Algorithm PEGE is initialized with the parameters $C(a) = h \cdot a$, for a tuning parameter $h > 0$, $\alpha = 1$ and $\beta = 1$, and the online game is played over T rounds, we get the following bound on expected regret:

$$R(T) \leq \sum_{x \in \sigma} \Delta_x \left(\frac{\log T}{h} \right)^2 + \frac{4\sqrt{2\pi} e^2 R \Delta_{max} \beta_\sigma}{\Delta} e^{\frac{h^2(2R^2\beta_\sigma^2)}{\Delta^2}}. \tag{18}$$

Proof. Let total number of phases that the algorithm runs for be K. We relate K to total time T as (after substituting parameters $C(a) = h \cdot a$, $\alpha = 1$ and $\beta = 1$ in Algorithm PEGE):

$T = \sum_{i=1}^{K} |\sigma| i + \sum_{i=1}^{K} e^{hi} \geq e^{hK} \implies K \leq \dfrac{\log T}{h}$.

Exploration regret: Sine we are in distribution dependent setting now, expected exploration regret in each exploration phase is $\sum_{x \in \sigma} \Delta_x$. Hence, total expected exploration regret is upper bounded by:

$\sum_{i=1}^{K} (\sum_{x \in \sigma} \Delta_x) i = \sum_{x \in \sigma} \Delta_x \frac{K(K+1)}{2} \leq (\sum_{x \in \sigma} \Delta_x) \dfrac{\log^2 T}{h^2}$.

Exploitation regret: When a sub-optimal arm is picked in an exploitation round, the expected regret in that round is: $\leq \Delta_{max}$. Using Lemma 6 with $\beta = 1$, the total expected regret due to exploitation over K phases is upper bounded by:

$$\begin{aligned}\sum_{i=1}^{K} \underbrace{2e^2\Delta_{max}\, e^{hi-\frac{i(i+1)}{2}\frac{\Delta^2}{4R^2\beta_\sigma^2}}}_{\text{expected exploitation regret upper bound in phase i}} &\leq 2e^2\Delta_{max}\sum_{i=1}^{\infty} e^{hi-\frac{i(i+1)}{2}\frac{\Delta^2}{4R^2\beta_\sigma^2}} \\ &\leq 2e^2\Delta_{max}\int_{-\infty}^{\infty} e^{hy-\frac{y(y+1)}{2}\frac{\Delta^2}{4R^2\beta_\sigma^2}}\,dy\end{aligned} \tag{19}$$

The integral is the moment generating function (adjusting for normalization constant) of a gaussian random variable $Y \in \mathcal{N}(0, \frac{4R^2\beta_\sigma^2}{\Delta^2})$. Thus, the integral is $\mathbb{E}[e^{hY}] = e^{\frac{2h^2R^2\beta_\sigma^2}{\Delta^2}}$ and total expected regret due to exploitation is upper bounded by: $\dfrac{4e^2\Delta_{max}\sqrt{2\pi}R\beta_\sigma}{\Delta}e^{\frac{2h^2R^2\beta_\sigma^2}{\Delta^2}}$.

Summing over exploration and exploitation regrets completes the proof. □

7.2 Proof of Results in Section 4

The following theorem is about the version of PEGE that Algorithm 3 calls on line 8. It will be needed in the proof of Theorem 4.

Theorem 7. (Distribution Dependent Regret, version 2) *When Algorithm 1 is initialized with the parameters $C(a) = h \cdot a$, for a tuning parameter $0 < h < \frac{\Delta^2}{4R^2\beta_\sigma^2}$, $\alpha = 1$ and $\beta = 0$, and the online game is played over T rounds, we get the following bound on expected regret:*

$$R(T) \leq \sum_{x\in\sigma}\Delta_x\frac{\log T}{h} + \frac{2e^2\Delta_{max}}{\frac{\Delta^2}{4R^2\beta_\sigma^2} - h} \tag{20}$$

Note: Compared to Theorem 2, the regret bound has better dependence on T — $O(\log T)$ instead of $O(\log^2 T)$ — but it also has a disadvantage. If the tuning parameter h is incorrectly set, say $h \geq \frac{\Delta^2}{4R^2\beta_\sigma^2}$, then the bound does not even apply.

Proof. **Key Steps**:

Let total number of phases that the algorithm runs for be K. First: $T = \sum_{i=1}^{K}|\sigma| + \sum_{i=1}^{K} e^{hi} \geq e^{hK}$ $\implies K \leq \dfrac{\log T}{h}$.

Expected regret due to exploration: $\sum_{i=1}^{K}(\sum_{x\in\sigma}\Delta_x) = \sum_{x\in\sigma}\Delta_x K \leq (\sum_{x\in\sigma}\Delta_x)\dfrac{\log T}{h}$.

Expected regret due to exploitation: When a sub-optimal arm is picked, expected regret $\leq \Delta_{max}$. Using Lemma 6 with $\beta = 0$, and tuning parameter $h < \frac{\Delta^2}{4R^2\beta_\sigma^2}$, we get total expected regret due to exploitation

$$\begin{aligned}2e^2\Delta_{max}\sum_{i=1}^{K} e^{hi-i\frac{\Delta^2}{4R^2\beta_\sigma^2}} &\leq 2e^2\Delta_{max}\sum_{i=1}^{\infty} e^{hi-i\frac{\Delta^2}{4R^2\beta_\sigma^2}} \\ &= 2e^2\Delta_{max}\sum_{i=1}^{\infty} e^{-i(\frac{\Delta^2}{4R^2\beta_\sigma^2}-h)} \\ &\leq 2e^2\Delta_{max}\int_0^{\infty} e^{-y(\frac{\Delta^2}{4R^2\beta_\sigma^2}-h)}dy \\ &= \frac{2e^2\Delta_{max}}{\frac{\Delta^2}{4R^2\beta_\sigma^2}-h}\end{aligned} \tag{21}$$

□

7.2.1 Proof of Theorem 3

Proof. Note that Assumption 1 through Assumption 3 hold. Therefore, from Lemma 5, with $\beta = 0$ we get, with probability at least $1 - \delta_b$,

$$\forall x,\ |\bar{r}(x, \theta_p^*) - \bar{r}(x, \hat{\theta}(b)| \leq \sqrt{\frac{R^2 \beta_\sigma^2 \log(\frac{2e^2}{\delta_b})}{b}}$$

Let $\delta_b = \delta/2b^2$ which implies $\sum_{b \geq 1} \delta_b = \pi^2 \delta/12 < \delta$. Thus, setting $w(b) = \sqrt{\frac{R^2 \beta_\sigma^2 \log(\frac{4e^2 b^2}{\delta})}{b}}$, the event E defined as

$$\forall b \geq 1, \forall x \in \mathcal{X}, |\bar{r}(x, \hat{\theta}(b)) - \bar{r}(x, \theta_p^*)| \leq w(b). \tag{22}$$

holds with probability at least $1 - \delta$.

1. Note that $b \geq T_1(\delta)$ implies $8w(b) < \Delta$. This is because the latter has the form $e^{Lb} > Mb$ with $M = 2e/\delta$ and $L = \Delta^2/(128R^2\beta_\sigma^2)$. Setting $b \geq 2/L \log(2M/L)$ guarantees that $e^{Lb/2} \geq 2M/L$ which implies that $e^{Lb} \geq Mb$ since $e^{Lb/2} \geq 1 + Lb/2 \geq Lb/2$.

   If $8w(b) < \Delta$ then clearly $2w(b) < \Delta$. Let $x \neq x^*$ be arbitrary. We have the following chain of implications:

   $$\begin{aligned}
   & 2w(b) < \Delta \\
   \Rightarrow\ & 2w(b) < \bar{r}(x^*, \theta_p^*) - \bar{r}(x_-^*, \theta_p^*) && \text{(def. of } \Delta\text{)} \\
   \Rightarrow\ & 2w(b) < \bar{r}(x^*, \theta_p^*) - \bar{r}(x, \theta_p^*) && \text{(Assumption 4)} \\
   \Rightarrow\ & 0 < \bar{r}(x^*, \hat{\theta}(b)) - \bar{r}(x, \hat{\theta}(b)). && (\because E \text{ holds})
   \end{aligned}$$

   This means that the If condition on line 12 will evaluate to true and $\hat{x}(b)$ on line 13 will be set to x^*.

   We also have the following chain of implications:

   $$\begin{aligned}
   & 8w(b) < \Delta \\
   \Rightarrow\ & 8w(b) < \bar{r}(x^*, \theta_p^*) - \bar{r}(x_-^*, \theta_p^*) && \text{(def. of } \Delta\text{)} \\
   \Rightarrow\ & 8w(b) < \bar{r}(x^*, \theta_p^*) - \bar{r}(\hat{x}_-(b), \theta_p^*) && (\because \bar{r}(\hat{x}_-(b), \theta_p^*) \leq \bar{r}(x_-^*, \theta_p^*)) \\
   \Rightarrow\ & 8w(b) < \bar{r}(\hat{x}(b), \theta_p^*) - \bar{r}(\hat{x}_-(b), \theta_p^*) && (\because \hat{x}(b) = x^*) \\
   \Rightarrow\ & 6w(b) < \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b)). && (\because E \text{ holds})
   \end{aligned}$$

   This means that the If condition on line 15 will evaluate to true and the algorithm will stop and output an estimate $\hat{\Delta}$.

   Now suppose the algorithm stops and does not output "threshold exceeded" which means that the If conditions on line 12 and line 15 were both true at some episode b. Let $x \neq \hat{x}(b)$ be arbitrary. We have the following chain of implications:

   $$\begin{aligned}
   & 6w(b) < \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) && \text{(line 15)} \\
   \Rightarrow\ & 6w(b) < \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(x, \hat{\theta}(b)) && (\hat{x}(b) \text{ unique maximizer by line 12}) \\
   \Rightarrow\ & 4w(b) < \bar{r}(\hat{x}(b), \theta_p^*) - \bar{r}(x, \theta_p^*). && (\because E \text{ holds})
   \end{aligned}$$

   This means, along with Assumption 4, that $\hat{x}(b) = x^*$. We also have,

   $$\begin{aligned}
   & 6w(b) < \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) && \text{(line 15)} \\
   \Rightarrow\ & 6w(b) < \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(x_-^*, \hat{\theta}(b)) && (\because \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) \geq \bar{r}(x_-^*, \hat{\theta}(b))) \\
   \Rightarrow\ & 4w(b) < \bar{r}(\hat{x}(b), \theta_p^*) - \bar{r}(x_-^*, \theta_p^*) && (\because E \text{ holds}) \\
   \Rightarrow\ & 4w(b) < \bar{r}(x^*, \theta_p^*) - \bar{r}(x_-^*, \theta_p^*) && (\because \hat{x}(b) = x^*) \\
   \Rightarrow\ & 4w(b) < \Delta. && \text{(def. of } \Delta\text{)}
   \end{aligned}$$

Now we prove that the output $\hat{\Delta}$ lies in the right range. We have

$$\begin{aligned}
\hat{\Delta} &= \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) && \text{(line 16)} \\
&\geq \bar{r}(\hat{x}(b), \theta_p^*) - \bar{r}(\hat{x}_-(b), \theta_p^*) - 2w(b) && (\because E \text{ holds}) \\
&= \bar{r}(x^*, \theta_p^*) - \bar{r}(\hat{x}_-(b), \theta_p^*) - 2w(b) && (\because \hat{x}(b) = x^*) \\
&\geq \bar{r}(x^*, \theta_p^*) - \bar{r}(x_-^*, \theta_p^*) - 2w(b) && (\because \bar{r}(\hat{x}_-(b), \theta_p^*) \leq \bar{r}(x_-^*, \theta_p^*)) \\
&\geq \Delta - 2w(b) && \text{(def. of } \Delta) \\
&\geq \frac{\Delta}{2}. && (\because w(b) < \Delta/4)
\end{aligned}$$

Similarly,

$$\begin{aligned}
\hat{\Delta} &= \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) && \text{(line 16)} \\
&\leq \bar{r}(\hat{x}(b), \hat{\theta}(b)) - \bar{r}(x_-^*, \hat{\theta}(b)) && (\because \bar{r}(\hat{x}_-(b), \hat{\theta}(b)) \geq \bar{r}(x_-^*, \hat{\theta}(b))) \\
&= \bar{r}(x^*, \hat{\theta}(b)) - \bar{r}(x_-^*, \hat{\theta}(b)) && (\because \hat{x}(b) = x^*) \\
&\leq \bar{r}(x^*, \theta_p^*) - \bar{r}(x_-^*, \theta_p^*) + 2w(b) && (\because E \text{ holds}) \\
&\leq \Delta + 2w(b) && \text{(def. of } \Delta) \\
&\leq \frac{3\Delta}{2}. && (\because w(b) < \Delta/4)
\end{aligned}$$

2. In this case $T_0 \leq T_1(\delta)$ but it could still be that the algorithm stops not because the threshold is exceeded but because line 12 and line 15 were true at some episode b. Clearly $b < T_0$, otherwise we would have output "threshold exceeded" and not produced an estimate $\hat{\Delta}$. Under the event E, the previous part shows that if stopping occurs with an estimate $\hat{\Delta}$, it must be that $4w(b) < \Delta$, i.e.

$$4\sqrt{\frac{R^2\beta_\sigma^2 \log(\frac{4e^2b^2}{\delta})}{b}} < \Delta \quad \Rightarrow \quad b > \frac{16R^2\beta_\sigma^2}{\Delta^2} \log \frac{4e^2}{\delta} = T_2(\delta).$$

   This means $T_0 > b > T_2(\delta)$.

3. Finally, suppose Assumptions 1 through 3 hold but Assumption 4 fails. Event E still holds with probability at least $1 - \delta$. However, if there are at least two optimal actions then, under E, their confidence intervals will always overlap and If condition on line 15 will never be true. That means that the algorithm can only stop when the threshold T_0 is exceeded.

□

7.2.2 Proof of Theorem 4

Proof. We break the proof into the two cases mentioned in the theorem statement.

Part 1: Assumption 4 fails. From Theorem 3 we know, that with probability at least $1 - \delta$, Algorithm 2 outputs "threshold exceeded" in this case. Because of Eq. (22), we also have, for an optimal action x^*:

$$|\bar{r}(\hat{x}(T_0), \theta_p^*) - \bar{r}(x^*, \theta_p^*)| \leq 2w(T_0)$$

which implies a total regret of

$$2w(T_0)(T - T_0|\sigma|) \leq 2w(T_0)T$$

in the remaining $T - T_0|\sigma|$ rounds since we execute line 5. The regret when Algorithm 2 was running is bounded by $R_{max}T_0|\sigma|$. On the bad event, which occurs with probability at most $1 - \delta$, the regret is at most TR_{max} giving us a total expected regret of

$$2w(T_0)T + T_0|\sigma|R_{max} + \delta TR_{max} = 2\sqrt{\frac{R^2\beta_\sigma^2 \log(\frac{4e^2T_0^2}{\delta})}{T_0}}T + T_0|\sigma|R_{max} + \delta TR_{max}$$

which is upper bounded by

$$2(2R\beta_\sigma|\sigma|^2R_{max}^2T)^{2/3}\sqrt{\log(4e^2T^3)}+R_{max}$$

for $T_0=\left(\frac{2R\beta_\sigma T}{|\sigma|R_{max}}\right)^{2/3}$ and $\delta=1/T$.

Part 2: Assumption 4 holds. *Case A:* $T_1(\delta)<T_0$. In this case, according to Theorem 3, with probability at least $1-\delta$, Algorithm 2 finishes in $T_1(\delta)$ episodes and outputs $0.5\Delta\leq\Delta\leq1.5\Delta$. This means $\Delta/36R^2\beta_\sigma^2\leq h\leq\Delta^2/4R^2\beta_\sigma^2$. Therefore, by Theorem 7, we have, regret due to Algorithm 1 is at most:

$$\sum_{x\in\sigma}\Delta_x\frac{36R^2\beta_\sigma^2\log T}{\Delta^2}+\frac{8e^2R^2\beta_\sigma^2}{\Delta^2}.$$

Overall, the expected regret is bounded by

$$T_1(\delta)R_{max}|\sigma|+\sum_{x\in\sigma}\Delta_x\frac{36R^2\beta_\sigma^2\log T}{\Delta^2}+\frac{8e^2R^2\beta_\sigma^2}{\Delta^2}+R_{max}T\delta.$$

For $\delta=1/T$, this becomes

$$\frac{256R^2\beta_\sigma^2}{\Delta^2}\log\frac{512e^2R^2\beta_\sigma^2T}{\Delta^2}R_{max}|\sigma|+\sum_{x\in\sigma}\Delta_x\frac{36R^2\beta_\sigma^2\log T}{\Delta^2}+\frac{8e^2R^2\beta_\sigma^2}{\Delta^2}+R_{max}.$$

Case B: $T_2(\delta)\leq T_0\leq T_1(\delta)$. In this regime, Algorithm 2 can stop and output "threshold exceeded", in which case, expected regret is bounded, as in Part 1, by

$$2(2R\beta_\sigma|\sigma|^2R_{max}^2T)^{2/3}\sqrt{\log(4e^2T^3)}+R_{max}.$$

However, it can also happen that Algorithm 2 stops and outputs $0.5\Delta\leq\Delta\leq1.5\Delta$ with probability at least $1-\delta$. In that case, total expected regret is bounded, as in Part 2, Case A, by

$$\frac{256R^2\beta_\sigma^2}{\Delta^2}\log\frac{512e^2R^2\beta_\sigma^2T}{\Delta^2}R_{max}|\sigma|+\sum_{x\in\sigma}\Delta_x\frac{36R^2\beta_\sigma^2\log T}{\Delta^2}+\frac{8e^2R^2\beta_\sigma^2}{\Delta^2}+R_{max}.$$

Note that the above bound scales as $O(T_1(\delta))$ for $\delta=1/T$, which is upper bounded by $O(T_2(\delta)\log T_2(\delta))$. But we know that $T_2(\delta)\leq T_0$ which means the bound is no larger than $O(T_0\log T_0)=O(T^{2/3}\log T)$. So no matter what happens, regret is upper bounded by $O(T^{2/3}\log T)$ in this case.

Case C: $T_0\leq T_2(\delta)$. With probability at least $1-\delta$, by Theorem 3, in this case, Algorithm cannot stop and output $\hat{\Delta}$. Instead, it outputs "threshold exceeded". When this happens, Algorithm 1 never gets called and only exploitation rounds follow (line 5). Regret is bounded, just as in Part 1, by

$$2(2R\beta_\sigma|\sigma|^2R_{max}^2T)^{2/3}\sqrt{\log(4e^2T^3)}+R_{max}.$$

□